

EARNEST: Experimental Analysis of RAN Energy with open-source Software Tools

Venkateswarlu Gudepu[§], Bhargav Chirumamilla[§], Rajashekhar Reddy Tella[§], Abhishek Bhattacharyya[§], Shubh Agarwal[§], Lavanya Malakalapalli[§], Carlo Centofanti[•], José Santos[†], Koteswararao Kondepu[§]

[§]Indian Institute of Technology Dharwad, Dharwad, India

[•]University of L'Aquila, DISIM Department, L'Aquila, Italy

[†] Ghent University - imec, IDLab, Department of Information Technology, Gent, Belgium

e-mail: 212011003@iitdh.ac.in, carlo.centofanti1@univaq.it, k.kondepu@iitdh.ac.in

Abstract—The 5G and Beyond (B5G) networks aim to support diverse use cases — extremely low latency, high data rates, and dense user connectivity. However, meeting these use cases results in an increase in energy consumption due to the use of computing resources in the B5G networks. While several studies focus on 5G Core Network (5G CN) energy consumption, it's essential to acknowledge that a substantial 75% of the network's overall energy usage occurs within the Radio Access Network (RAN). Hence, it's crucial to focus on RAN energy performance improvements. This paper exploits various open-source software tools that measure and monitor RAN energy, which helps design energy-efficient RAN. Different RAN architectures such as Monolithic, Disaggregated, and Control Plane and User Plane Separation (CUPS) are considered to measure and monitor energy consumption using open-source software tools — *S-tui* and *Scaphandre*. We study energy consumption as a function of the number of connected User Equipments (UEs) and the impact of connecting multiple Distributed Units (DUs) on the energy consumption of both the Control Plane (CP) and User Plane (UP) of gNB Central Unit (CU). We also study the energy consumption of various open source 5G CN. Finally, this study examines the influence of various RAN parameters on energy consumption by using a real-time dataset of the monolithic RAN scenario.

Index Terms—Radio Access Network (RAN), Energy Efficiency, Power consumption, 5G and Beyond (B5G).

I. INTRODUCTION

The United Nations (UN) General Assembly introduced seventeen interconnected Sustainable Development Goals (SDGs) to shape “a better and more sustainable future for all” by 2030 [1]. Among the many industries, Information and Communication Technologies (ICT), including wireless networks, have been identified as crucial in contributing to all the seventeen SDGs. In this context, ongoing research on the 5G and Beyond (B5G) networks is aligned with these global goals, with the aim of using their architectures to actively support the realization of the SDGs, including the critical goals of reducing both the energy consumption and carbon footprint [2]. The ICT sector's sustainability commitment necessitates solutions for measuring energy consumption, carbon emissions, and enhancing network performance. Recent studies indicate that in 2020, the ICT sector contributes to global greenhouse gas emissions from 1.8% to 2.8%, with a substantial increase in energy consumption [3].

With the introduction of B5G networks and an increased variety of network services with data-hungry applications,

network operations have become increasingly complex. Also, the RAN is coping with the diverse use cases and devices while meeting the stringent Quality of Service (QoS) demands. This could result in an increase in energy consumption due to the exponential increase in data traffic, the large number of connected devices, and the use of computing resources in the network. Moreover, [4] describes that RAN contributes more than 75% of the energy consumption of the service provider network.

To effectively manage RAN energy, there is a strong need for software/hardware-based tools to measure and monitor RAN energy consumption. An increase in power consumption signifies that the network is drawing more energy from its resources. Energy monitoring empowers operators to track and reduce carbon emissions, adhere to regulatory requirements, and implement practices that promote energy efficiency. Beyond these benefits, it provides valuable insights into capacity planning, performance optimization, and early issue detection, supporting operational excellence and the long-term sustainability objectives inherent to the dynamic nature of B5G services.

A work in [5], measures the power consumption during the user registration and authentication of various deployed open source 5G CN such as free5GC, open5GS, and OpenAirInterface (OAI). Other work in [6], uses a hardware tool (i.e., Meross MSS310) to measure the power consumption at the 5G CN. However, these works are focusing on the power measurement at the 5G CN only. Another work in [7], presents an optimization method for power consumption of the RAN, whereas [4], measures the energy consumption of a RAN component — Radio Unit (RU) or Base Band Unit (BBU) processing — overall the network energy usage. Both these works focus on measuring the energy using analytical methods.

However, an experimental investigation on energy measurement and monitoring at the RAN is essential to design the energy efficient RAN for B5G networks. The main contributions of this paper are as follows:

- Exploring various open source tools to understand the energy/power consumption of the RAN — Monolithic, Disaggregated, and CUPS.
- Investigation on the power consumption as a function of

- a number of connected CUPS, DUs, and UEs.
- Evaluating the power consumption of various open source 5G CNs — OAI5G-CN and Open5GS.
- Analyzing the impact of various RAN parameters as a function of power consumption.

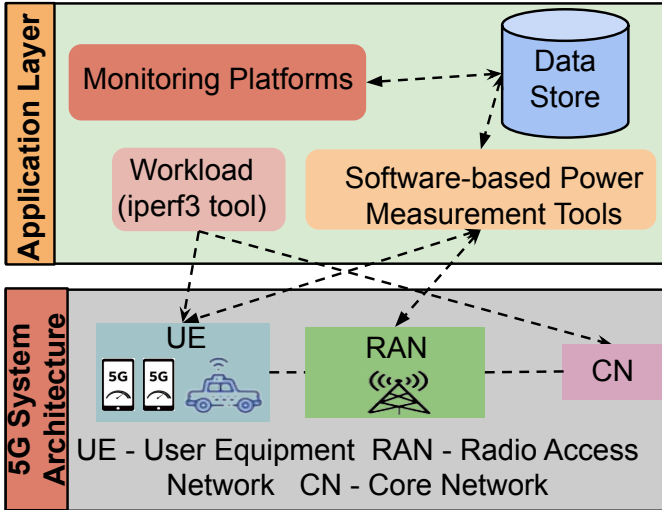


Fig. 1: System Model

II. SYSTEM MODEL

This section describes the system model, which includes two components — 5G system architecture and the application layer. The 5G system architecture consists of three key components — UEs, RAN, and 5G CN as shown in Fig. 1. These components work together to enable the functionality and capabilities of 5G networks. The UE serves as the endpoint device for network connectivity, the RAN facilitates wireless communication between the UE and the network, and the 5G CN manages data routing, network management, and various services. The RAN part is evolving over time as mainly three different deployments (i) Monolithic RAN — a single gNB; (ii) Disaggregated RAN — where gNB is decoupled into DUs and Central Units (CUs), and (iii) CUPS RAN — where the CU is further decoupled into Control Plane (CP) and User Plane (UP) — to meet the diverse demands of 5G networks.

Whereas the application layer in Fig. 1 includes open-source software-based power measurement tools, which are instrumental in measuring power consumption across different deployments of 5G system architecture, offering insights into power consumption and carbon emission. The data store serves as a repository for extracting and storing power consumption data, enabling in-depth analysis using various monitoring platforms. The monitoring platforms provide real-time visibility into power consumption metrics and system performance. An iperf3 tool [8] is utilized to generate workloads that assess the end-to-end performance of the 5G system, which helps to understand its power usage and efficiency.

The components within the application layer interact with the end-to-end running 5G system architecture. A software-based power measurement tool measures power consumption

across the RAN and UE, storing this data in a database. The extracted metrics are then utilized by monitoring platforms to gain deeper insights into the power consumption of 5G system components.

III. EXPERIMENTAL SETUP

This section describes the experimental setup used to measure and monitor the power consumption of the 5G system architecture. We conduct the experiments for three scenarios — Monolithic, Disaggregated, and CUPS RAN as shown in Fig. 2 to demonstrate the power consumption of the RAN and description of each scenario as follows:

A. Scenario 1: Monolithic RAN

It combines all RAN functions, such as the Physical layer, Medium Access Control (MAC) layer, Radio Link Control (RLC) layer into a single, integrated unit called Next-generation NodeB (gNB). The monolithic RAN setup (i.e., scenario 1) contains gNB, UE, and 5G CN. Fig. 2 shows the experimental setup used to measure the power consumption on the monolithic RAN. Each component of the monolithic RAN is deployed on separate host machines with the system configuration as shown in Table I. Despite having advantages like easier deployment since all the components are tightly coupled, it poses certain challenges like lack of scalability, Vendor Lock-In, performance issues, slow speed of development, etc. These challenges are resolved by the disaggregated RAN scenario.

B. Scenario 2: Disaggregated RAN

In this scenario, the gNB functions are physically decoupled into RU, DU, and CU by providing various split options [9]. The advantages of using the disaggregated architecture are high flexibility, scalability, resource optimization, among others. Fig. 2 shows the experimental setup used for measuring the power consumption over the disaggregated RAN.

Each component of the disaggregated RAN — CU, DU, and UE is deployed on separate host machines with the system configuration as shown in Table I. However, to optimize the RAN functions are further disaggregated the gNB-CU into gNB-CU-Control Plane (gNB-CU-CP) and gNB-CU-User Plane (gNB-CU-UP) to support different use case scenarios.

C. Scenario 3: CUPS RAN

In this architecture, the gNB-CU further splits into gNB-CU-CP and gNB-CU-UP. The CP is responsible for the initial attachment procedure between the RAN components and the UP is responsible for forwarding the data packets. This kind of separation brings several advantages in the form of better scalability of UP operations, such as the ability to design user planes, the ability to maintain and upgrade systems, etc. It also reduces the operational and maintenance costs of the Mobile Network Operators (MNOs) by enabling the hosting of the CP and UP nodes at different geographical locations. The experimental setup used for measuring and monitoring the power consumption over the CUPS RAN is as shown in Fig. 2. Each component of the CUPS RAN — gNB-CU-CP, gNB-CU-UP, DUs, and UE — is deployed on separate host machines with the system configuration as shown in Table I.

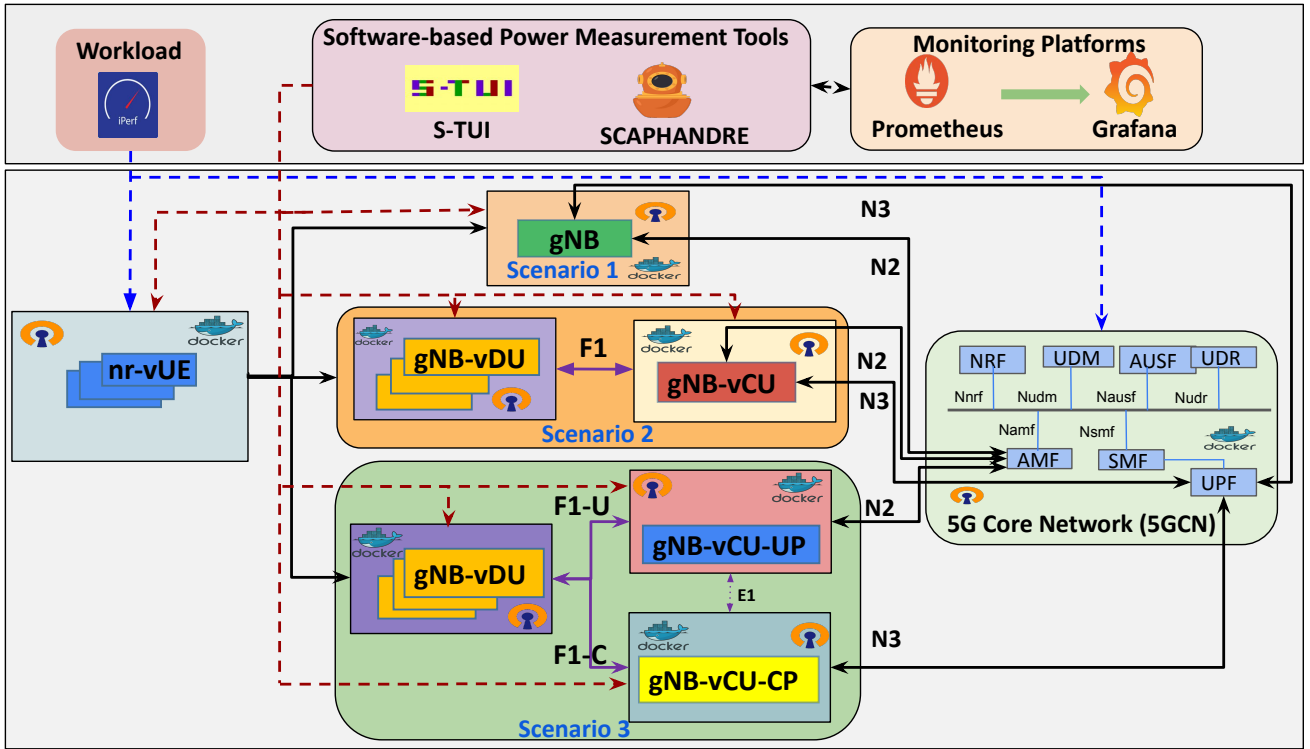


Fig. 2: Experimental setup for different scenarios

In all three scenarios, we used OAI software packages to build all the RAN and 5G CN components. The radio functionality is implemented using an RF simulator but can also be extended by using Software Defined Radio (SDR) boards such as NI B210 or Universal Software Radio Peripheral (USRP) N310. The OAI develop branch is used for deploying the RAN units — the UE, DU, and CU — and the version for the 5G CN is v1.2.1. The OAI 5G New Radio (5G-NR) provides different split options as defined in [9] and in our case, we make use of option 2 level split between gNB-DU and gNB-CU in case of our disaggregated setup and the experimental testbed is as shown in Fig. 3.

The CUPS setup contains gNB-CU-CP, gNB-CU-UP, gNB-DU, UE and 5G CN. Here the E1 interface is used to connect gNB-CU-CP and gNB-CU-UP and the F1 interface to connect gNB-CU and gNB-DU same as mentioned in the disaggregated one. The branch used for gNB-CU-CP, gNB-CU-UP, gNB-DU, and UE is *elap-implementation*. In this setup, a single gNB-CU-CP is connected to a single gNB-CU-UP, which in turn is connected to one or multiple gNB-DUs which are connected to respective UEs.

The implemented CN component in all three scenarios includes: Unified Data Repository (UDR), Unified Data Management (UDM), Authentication Server Function (AUSF), Network Repository Function (NRF), Access and Mobility Management Function (AMF), Session Management Function (SMF), User Plane Function (UPF) (i.e., SPGW-U). These functions are deployed as multiple docker containers using the OAI 5G CN and also deployed Open5GS to understand the 5G CN power consumption. The network parameters of

the deployed 5G system architecture are reported in Table II.

TABLE I: 5G System Components Configuration

Component	CPU	Memory	RAM
UE	11th Gen Intel(R) Core(TM) i7-11700 @ 2.50GHz	1 TB	32 GB
gNB	Intel(R) Xeon(R) W-2265 CPU @ 3.50GHz	1 TB	64 GB
DU	Intel(R) Xeon(R) W-2265 CPU @ 3.50GHz	2 TB	64 GB
CU	Intel(R) Core(TM) i5-7500 CPU @ 3.40GHz	1 TB	8 GB
CU-CP	Intel(R) Core(TM) i5-7500 CPU @ 3.40GHz	1 TB	8 GB
CU-UP	Intel(R) Core(TM) i5-7500 CPU @ 3.40GHz	1 TB	8 GB
CN	Intel(R) Core(TM) i5-7500 CPU @ 3.40GHz	1 TB	8 GB

TABLE II: 5G Network Parameters

Description	Value
NR Release	3GPP Release 16
NR Band	Band 78
NR Frequency	3.6 GHz
RAN type	5G standalone gNB
CU/DU split	Option 2
Physical Resource Block (PRB)	106
Radio Channel Bandwidth	40 MHz
Midhaul Capacity	1 Gbps Ethernet
Backhaul Capacity	1 Gbps Ethernet
UE	OAI based 5G SA UE

In order to measure the energy, we used two software-based tools — *S-tui*, and *Scaphandre* — across all the scenarios. *S-tui* provides real-time system performance monitoring and stress-testing in command-line environments, while the *Scaphandre*

measures power consumption using hardware sensors, correlating it with resource utilization and process activities, facilitating comprehensive power analysis. The collected power metrics are monitored using both Prometheus and Grafana. Prometheus is a widely used open-source monitoring tool that is used to scrape the system resource metrics from various exporters, and it also provides a powerful query language, i.e., PromQL. It enables the creation of custom alerting rules and provides flexible querying capabilities for in-depth analysis. On the other hand, Grafana complements Prometheus with its user-friendly visualization platform, allowing the creation of interactive and visually appealing dashboards that present power consumption data in a comprehensible manner. This integration facilitates effective monitoring, analysis, and optimization of power usage within various systems and infrastructures. Also, we employed the iperf3 tool to create the workload (i.e., uplink throughput of 15 Mbps between the UE and the 5G CN) and observe the power consumption throughout all the scenarios.

We evaluated the performance of two open source software tools — *S-tui* and *Scaphandre* for all three scenarios to measure the power consumption of the host machine, in which each component of the RAN is deployed.

IV. EXPERIMENTAL RESULTS

This section describes three different experimental results: (i) OAI deployment scenarios results; and (ii) Core Network deployment results.

A. Deployment scenarios results

The experiments of above described scenarios are conducted over a period of 240 *sec*, with a workload generated using the iperf3 tool within the [30 – 210] *sec* interval. For each considered scenario, the *idle* scenario describes where none of the RAN components is running — represented with *gNB idle*, *UE idle*, *DU idle*, *CU idle*, *CU-CP idle*, *CU-UP idle* — this shows the base power consumed by the corresponding systems. The power consumption of the monolithic RAN (i.e., scenario 1) as a function of experiment time is shown in both Fig. 4a and 4b, measured by the *S-tui* and *Scaphandre* tools, respectively. When the gNB initially started, the observed power consumption is around 40W and increased to around 60W during the workload — experiment time of first 30 secs — which can be attributed to the high data traffic and the required system resource allocation. Interestingly, the *S-tui* and *Scaphandre* both the considered tools are able to provide approximately similar behavior. This could validate that the experimental results are consistent.

The power consumption of the disaggregated RAN (i.e., scenario 2) as a function of experiment time is depicted in both Fig. 4c and 4d, where power consumption is measured by the *S-tui* and *Scaphandre* tools, respectively. When the CU and DU initially started, the observed power consumption was 15W and 37W, respectively. During the workload, a slight increase of 2.67% on the CU side and an increase of 11.9% can be observed against the initial power consumption. This

increase can be attributed to the functionalities distributed between both CU and DU within the disaggregated architecture. The experimental results confirm that the disaggregated RAN could provide lower energy consumption when compared to monolithic gNB by splitting monolithic RAN into two different components.

Furthermore, in Fig. 5, the power consumption is measured as a function of the number of UEs, where all the UEs are deployed as containers due to the lack of physical machines. Note that there is a slight increase in power consumption compared to Fig. 4c; this is due to changing UE deployment from bare metal to containerization. The increase in the number of UEs resulted in a 19.3% increase in the power consumption. Here, only 1 UE is performing the workload, and the remaining UEs are in a connection state.

The power consumption of the CUPS RAN (i.e., scenario 3) as a function of experiment time is depicted in both Fig. 4e and 4f, where power consumption is measured by the *S-tui* and *Scaphandre* tools, respectively. When the CU-CP, CU-UP, and DU started, the observed power consumption is 16.7W, 15.7W, and 38W, respectively. During the workload, the power consumption of CU-CP, CU-UP, and DU increased by 16.8%, 16.2% and 44%, compared to their initial power consumption. The slight increase in power consumption of both CU-CP and CU-UP, but a significant increase in DU can be attributed to the creation of routing tables, packet forwarding and processing of various layers, respectively. Additionally, Fig. 6 illustrates the power consumption of CU-CP and CU-UP as a function of the number of DUs. Fig. 6 shows that both CU-CP and CU-UP experience a 54% and 52% increase in power consumption, respectively, when four DUs connected, compared to a single connected DU. This can be attributed to the active exchange of the messages among the CUPS RAN components whenever a DU and corresponding UE are connected.

B. 5G CN energy experiment results

We present the power consumption of two 5G CN open sources — OAI and Open5GS while running the end-to-end connected 5G setup as shown in Fig 7. When the Open5GS and OAI5G-CN are initially started, the observed power consumption is around 11W and 34.7W, respectively. During the workload, the Open5GS experiences an 18% increase, while OAI5G-CN shows a 7.6% increase in power compared to their initial power consumption. Note that the OAI5G-CN power consumption increases during the workload load and decreases afterward, whereas the Open5GS power consumption does not reduce after the workload, which is specific to its implementation. However, OAI5G-CN consumes more power compared to the Open5GS. This could be attributed to the fact that in Open5GS all the 5G CN functions run in a bare metal. However in the OAI 5G CN implementation all the 5G CN functions are running in the form of Virtualized Network Functions (VNFs). These VNFs are deployed as multiple docker containers running through a docker compose YAML file. This containerized deployment of the 5G CN functions is one the reasons there is additional energy consumption.



Fig. 3: Disaggregated RAN experimental testbed

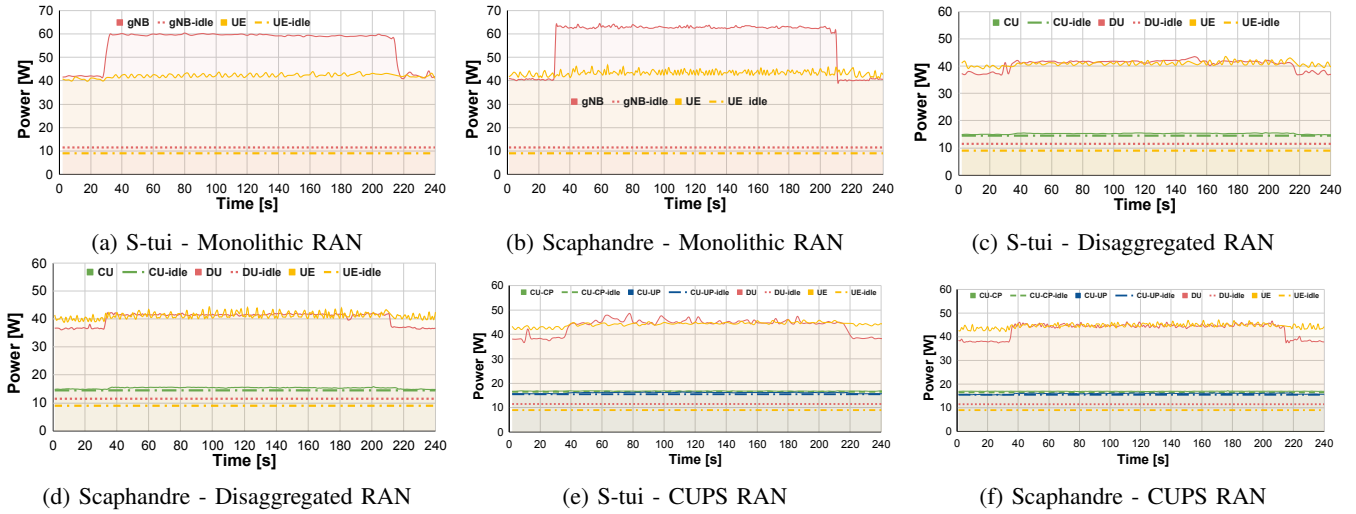


Fig. 4: Power consumption measurement using *S-tui*, and *Scaphandre* for different RAN scenarios.

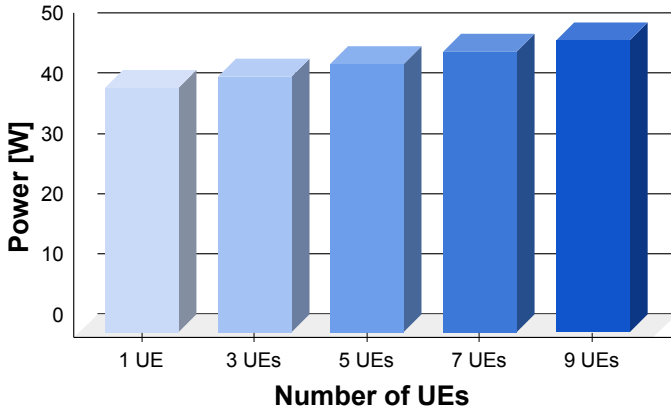


Fig. 5: Power consumption as a function of number of UEs.

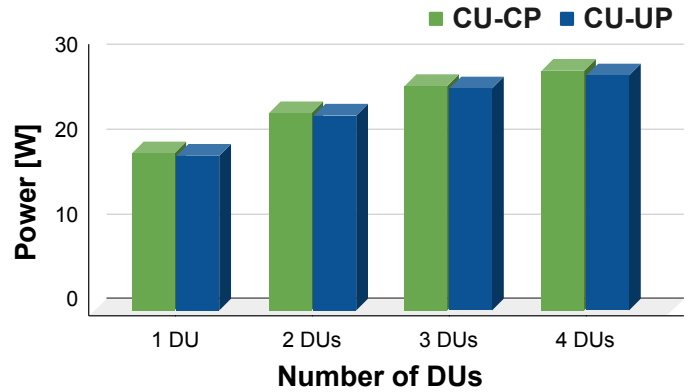


Fig. 6: Power consumption of both CU-CP and CU-UP.

V. IMPACT OF RAN PARAMETERS ON ENERGY CONSUMPTION

In addition to power measurement and monitoring, an investigation into various parameters of the RAN is conducted to understand their impact on power consumption. This comprehensive analysis aims to identify specific parameters that can be fine-tuned to enhance RAN energy efficiency further. We considered a real-time dataset [10], which is collected from a monolithic gNB over the Open-RAN testbed [11]. The collected gNB parameters were taken under two different scenarios — (i) the uplink channel and (ii) downlink and uplink channels. It consists of several gNB parameters —

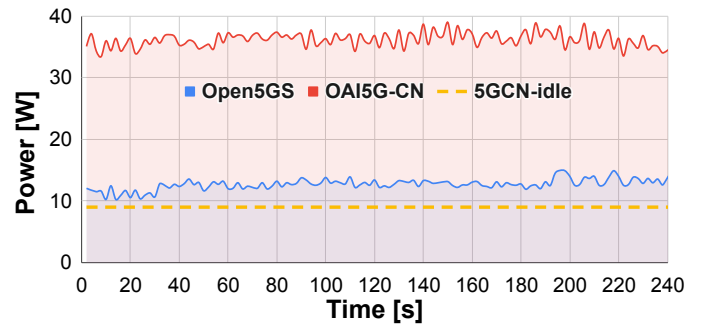


Fig. 7: Power consumption of Open5GS and OAI5G-CN.




Fig. 8: SHAP analysis for RAN parameters influence on power consumption

e.g., average power, SNR, MCS, subframe decoding time. To gain insights into the factors affecting power consumption, we leveraged SHapley Additive exPlanations (SHAP) analysis [12]. It is a powerful technique used to explain the output of Machine Learning (ML) models by quantifying the contribution of each input feature to the model's predictions. A Random Forest regressor has been applied to predict the average power and consider all the other gNB parameters as features of the AI/ML model [13].

Fig. 8 depicts the impact of various RAN parameters on the average power for two different scenarios. In Fig. 8a and 8b, the parameters are sorted by their magnitude of SHAP values. Red dots indicate a high impact, while blue dots signify a low impact on the average power. In Fig. 8a, the average subframe decoding time exhibits the highest magnitude of SHAP value, making it the most influential factor. Its negative magnitude implies that increasing this parameter results in a decrease in average power. The other parameters, such as average block error rate, average number of resource blocks, and average throughput, also have a substantial impact on average power. And in Fig. 8b, the parameter with the most impact on average power is the average throughput uplink, followed by airtime uplink, airtime downlink, and average throughput downlink. All these parameters have positive magnitude and significantly affect the average power.

VI. CONCLUSIONS AND FUTURE WORK

This paper addresses energy consumption in B5G networks, focusing on RAN. Using open-source software tools like *S-tui* and *Scaphandre*, we measured RAN power consumption in Monolithic, Disaggregated, and CUPS RAN scenarios. Our findings highlight energy performance insights, revealing a 19.3% increase in power consumption (Disaggregated RAN) with UEs ranging from 1 to 10. In CUPS RAN, CU-CP and CU-UP power consumption increased by 54% and 52%, respectively, with an increase in DUs. Evaluation of Open5GS and OAI5G-CN indicates OAI5G-CN consumes more power than Open5GS. The study also explores RAN parameter impact on energy consumption in real-time monolithic RAN dataset, offering a comprehensive understanding of factors influencing RAN energy use.

Based on this study, the potential future directions of this work are as follows: (i) *Exploring other hardware and software tools*: involves the exploration of additional hardware and

software solutions for RAN energy measurement and monitoring, enabling a more comprehensive assessment of energy consumption. (ii) *Optimizing the RAN parameters*: includes fine-tuning specific parameters within different RAN scenarios to achieve significant energy savings. (iii) *Energy estimation*: the exploration of AI/ML and predictive analytics for real-time energy consumption forecasting.

ACKNOWLEDGMENT

This work received funding from DST SERB Startup Research Grant (SRG-2021-001522), the SGNF project ("Reliability Evaluation of Virtualised 5G"). José Santos is funded by the Research Foundation Flanders (FWO), grant number 1299323N.

REFERENCES

- [1] V. Gudepu, B. Pappu, T. Javvadi *et al.*, "Edge computing in micro data centers for firefighting in residential areas of future smart cities," in *2022 International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*. IEEE, 2022, pp. 1–6.
- [2] J. Lorincz, A. Capone, and J. Wu, "Greener, energy-efficient and sustainable networks: State-of-the-art and new trends," p. 4864, 2019.
- [3] C. Freitag, M. Berners-Lee, K. Widdicks *et al.*, "The real climate and transformative impact of ict: A critique of estimates, trends, and regulations," *Patterns*, vol. 2, no. 9, 2021.
- [4] L. M. Larsen, H. L. Christiansen, S. Ruepp *et al.*, "Toward greener 5g and beyond radio access networks—a survey," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 768–797, 2023.
- [5] G. Lando, L. A. F. Schierholt, M. P. Milesi, and J. A. Wickboldt, "Evaluating the performance of open source software implementations of the 5g network core," in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2023, pp. 1–7.
- [6] A. Bellin, M. Centenaro, and F. Granelli, "A preliminary study on the power consumption of virtualized edge 5g core networks," in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*. IEEE, 2023, pp. 420–425.
- [7] A. Abrol and R. K. Jha, "Power optimization in 5g networks: A step towards green communication," *IEEE Access*, vol. 4, pp. 1355–1374, 2016.
- [8] M. Mortimer, "iperf3 documentation," 2018.
- [9] ITU-T, "Transport network support of IMT-2020/5G," ITU-T Technical Report, Feb. 2018.
- [10] J. X. Salvat Lozano, J. A. Ayala-Romero, L. Zanzi *et al.*, "O-ran experimental evaluation datasets," 2022. [Online]. Available: <https://dx.doi.org/10.21227/64s5-q431>
- [11] V. Gudepu, V. R. Chintapalli, P. Castoldi *et al.*, "Adaptive retraining of ai/ml model for beyond 5g networks: A predictive approach," in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*. IEEE, 2023, pp. 282–286.
- [12] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in neural information processing systems*, vol. 30, 2017.
- [13] V. R. Chintapalli, V. Gudepu, K. Kondepudi *et al.*, "Wip: Impact of ai/ml model adaptation on ran control loop response time," in *2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. IEEE, 2022, pp. 181–184.